Private Information and Price Regulation
In the US Credit Card Market

By

Scott Thomas Nelson

B.A. Economics and Mathematics
Yale University, 2010

SUBMITTED TO THE DEPARTMENT OF ECONOMICS IN PARTIAL
FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF

MASTER OF SCIENCE IN ECONOMICS
AT THE
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

FEBRUARY 2017

Signature of Author: .... Signature redacted

Department of Economics
January 20, 2017

Certified by: .... Signature redacted

James M. Poterba
Mitsui Professor of Economics
Thesis Supervisor

Accepted by: .... Signature redacted ....

Ricardo Caballero
~~Ford~~ International Professor of Economics
Chairman, Departmental Committee on Graduate Studies

Private Information and Price Regulation In the US Credit Card Market

By

Scott Thomas Nelson

Submitted to the Department of Economics
on January 20, 2017 in Partial Fulfillment of the
Requirements for the Degree of Master of Science in
Economics

ABSTRACT

Lenders typically learn new information about their borrowers over time but can be restricted from re-pricing debt in response to this information. I study a leading example of such re-pricing restrictions, the 2009 Credit CARD Act, to ask how such restrictions affect credit market efficiency. Using a near-universe of US consumer credit card account data as well as a large random sample of US consumer credit reports, I show evidence that the Act's restrictions had two competing effects: on the one hand, a decoupling between prices and default risk on existing loans over time, which engenders adverse selection through higher attrition of safe borrowers; on the other hand, lower markups on borrowers revealed to be inelastic, and hence lower price dispersion in the market overall. To quantify these two forces' net effect on market efficiency, I build a model of a competitive credit market with private information and changing borrower types over time, and I use the model to ask whether, and for whom, the Act's restrictions bring prices closer to an efficient benchmark of prices equaling marginal costs. While fully estimating the model remains a goal for future work, I here show preliminary results of how the model estimation is proceeding.

Thesis Supervisor: James M. Poterba
Title: Mitsui Professor of Economics

ACKNOWLEDGEMENTS AND DISCLAIMER

The views expressed herein are those of the author and do not necessarily reflect those of the Consumer Financial Protection Bureau or the United States. For their generous help at all stages of this project, I am deeply indebted to my dissertation committee: Antoinette Schoar, Jonathan Parker, and especially my committee chair Jim Poterba. For thoughtful discussions and suggestions, I thank Alexei Alexandrov, Alex Bartik, Vivek Bhattacharya, Ron Borzekowsi, Ken Brevoort, Tom Conkling, Amy Finkelstein, Daniel Green, Daniel Grodzicki, Greg Howard, Vikram Jambulapati, Luu Nguyen, Paul Rothstein, Richard Schmalensee, David Silberman, Daniel Waldinger, Jialan Wang, Mike Whinston, Wei Zhang, and David Zimmerman. I am also deeply indebted to Shaista Ahmed, Julian Jamison, Jesse Leary, Joe Remy, and Stefano Sciolli for their support and guidance at the CFPB. I gratefully acknowledge financial support from the National Science Foundation Graduate Research Fellowship under Grant No. 122374 and the MIT Bradley Fellowship.

I. Introduction

A typical prescription to address credit market information asymmetries is to generate more information: credit registries, private screening technologies, and courts all reveal information to lenders that otherwise would be private to borrowers. In this paper I challenge the conventional wisdom that such information revelation – and lenders' pricing on it – necessarily makes credit markets more efficient. When lenders have market power, the question of whether it is socially optimal for lenders to price on a piece of information depends on the mix of demand uncertainty and risk uncertainty that is resolved by that information. When demand uncertainty dominates this mix, the efficiency costs of pricing on such information can indeed overwhelm the benefits.

Empirically, I examine this question by evaluating a leading example of public policy that restricts what information lenders can use to price loans, the 2009 Credit CARD Act. Studying the CARD Act is also of independent interest, as it is arguably the most substantial regulation to affect the US consumer credit card market to date. Whereas a difference-in-differences analysis suggests the Act lowered the total cost of credit card borrowing by roughly 12 billion dollars per year (Agarwal et al., 2015), substantial questions remain about the Act's effects. In particular, if the Act did cause lower borrowing costs, it is unknown whether these price changes were the result of lower rents for lenders, or rather the mispricing of risk. Additionally, little is known about how the Act affected borrower exit from, or entry into, the market, and how this changing composition of borrowers contributed to the changes in loan pricing brought about by the Act.

A priori there is reason to believe the Act could both reduce lenders' rents and induce the mispricing of risk; both effects in turn can change the composition of individuals who choose to borrow. The CARD Act restricted consumer credit card lenders from changing borrowers' interest rates on outstanding debt in response to changes in borrowers' credit scores, repayment behavior, other borrowing, and nearly any other behavior observed either publicly or privately.[1] The Act also restricted a close substitute of interest rate repricing, namely, behavior-contingent fees. In practice, the incidence of upward repricing of interest on existing debt fell from over 10% of accounts on average per quarter in 2008-2009, to nearly 0% after the Act took effect, while total fee revenue also fell after the Act on the order of several billion dollars per year (CFPB 2013). To the extent that interest-rate repricing and the application of these fees had been driven by lenders' private information about borrower demand, these restrictions would be expected a priori to reduce rents; conversely, to the extent repricing had been in response to new or newly revealed borrower risk, these restrictions would be expected to induce risk mispricing.

I begin my analysis by presenting evidence that the Act both lowered lender rents and also exacerbated the mispricing of risk. I show, first, that upward repricing had been used to extract rents, in the sense that it targeted privately observed borrower behaviors (i.e. behaviors not reported to credit bureaus) that demonstrably were not predictive of risk and plausibly were predictive of price sensitivity. I confirm that the Act effectively eliminated such repricing. I then show, conversely, that the Act's price restrictions exacerbated the mispricing of risk. I show both that lenders possessed and priced for private information about risk before the Act, and that lender pricing became almost completely unresponsive to new information about risk after the Act. Connecting these changes to borrowers' participation choices, I show that borrowers whose riskiness had risen, and who therefore directly benefited from the Act's repricing restrictions, became less likely to attrite from their current lender after the Act. This adverse selection (via retention) of risky borrowers, together with the

[1] The sole exception to this restriction was in cases of borrowers paying late by 60 days or more to the lender in question, an exception which few lenders took advantage of in light of already-low recovery rates on 60-day past-due loans. Repricing was also allowed in some cases that were *not* dependent on borrower behavior, such as the movement of an index rate, or the expiration of a promotional interest rate on a pre-disclosed date not determined by borrower behavior.

preceding evidence of reduced lender rents, illustrates an example of the key efficiency tradeoff that emerges from the Act's pricing restrictions.

I then develop a quantitative model of the credit card market to examine how these two effects of the Act's pricing restrictions – reduced rents and the mispricing of risk – interact in equilibrium to affect borrowers' prices of borrowing and choices to borrow. The main goal of this modeling exercise is to assess whether, and for whom, the Act's price restrictions brought equilibrium prices closer to an efficiency benchmark of prices equaling marginal costs. Relatedly, I assess for which types of borrowers the Act's price restrictions affected the net subsidies paid from (or to) other more (or less) profitable borrowers in equilibrium. While fully estimating the model remains a goal for future work, I show preliminary results in this draft of how the model estimation is proceeding.

Central to estimating the model is a novel methodology to identify and measure borrowers' private information and its evolution over time. This methodology builds on tools from a large literature in health insurance that uses various ex-post outcomes to identify ex-ante risk types. I additionally exploit a novel source of price variation in consumer credit markets – occasional, idiosyncratic portfolio-wide repricing by certain lenders – in order to ground the model in well-identified demand estimates. This price variation obviates the usual need in similar models with imperfect competition to instrument for prices using, for example, competitors' product characteristics.

My work in this paper connects to a growing literature on public policy restricting pricing behavior in markets that have information problems. In the health insurance context, Handel et al. (2016) examine the effects of pricing restrictions similar to those in the Affordable Care Act, where insurers are prohibited from adjusting premia in response to (most) changes in health status over time. Their analysis focuses on the tradeoff between adverse selection and reclassification risk, as is also examined in Finkelstein et al. (2005) and Cochrane (1995). In contrast, I focus on the tradeoff between adverse selection (among other information problems) and lenders' rents, both because these rents have been salient in academic and policy debates about the credit card market, and because credit cards (like all unsecured credit) are arguably not designed for, and not cost-sustainably suited for, providing long-term insurance against permanent shocks. To my knowledge, this paper is the first to examine these questions in the context of a credit market, and to do so with a dataset representing the near-universe of the participants of the market in question.

This paper also relates to recent analyses of consumer switch costs in finance or insurance markets, which are an important source of lender market power in my analysis. Handel (2013), Illanes (2016), and Polyakova (2016) respectively study switch costs in health insurance, pensions, and prescription drug insurance. A long-standing literature in finance examines how borrower stickiness can either help mitigate information problems (Petersen and Rajan, 1995) or exacerbate pricing distortions (Sharpe 1990).

Finally, this paper also contributes to a large body of research on the credit card market in particular. The credit card market has been studied previously by, among others: Agarwal et al. (2010) and Ausubel (1999), who present evidence of adverse selection in the market; Ausubel (1991), Berlin and Mester (2004), Calem and Mester (1995), and Grodzicki (2014), who examine the competitiveness (or lack thereof) of credit card pricing; Drozd and Serrano-Padial (2014a, 2014b) who study refinancing risk and borrower screening technology; Ru and Schoar (2016) and Grodzicki (2015), who study the supply of direct-mail credit card offers; Santucci (2015), Jambulapati and Stavins (2014), and Agarwal et al. (2015), who respectively present evidence on lenders' supply of credit limits, account closures, and pricing changes through the implementation of the CARD Act; Keys and Wang (2015), who study the CARD Act's "nudges" for faster borrower repayment of debt; Gross and Souleles (2002) and

Agarwal et al. (2016), who examine borrowers' response to credit limit increases and lenders' corresponding incentives to lend; Brito and Hartley (1995), Stango and Zinman (2011), and Heidhues and Koszegi (2015), who study behavioral aspects of borrower behavior; Fulford (2014), Telyukova and Wright (2008), Gross and Souleles (2002), and Stango and Zinman (2016), who study credit card debt's role in household portfolio and consumption choices; and by Sullivan (2008) and Ganong and Noel (2016), who study credit card debt's role in insuring transitory employment shocks. There is also a small body of work focused on the CARD Act's repricing restrictions in particular: Hong et al. (2015) and Ronen and Pinheiro (2015) present theoretical models of the effects of repricing restrictions, while Levitin (2011) and Bar-Gill and Bubb (2011) present empirical results from the perspective of the law literature.

This long literature on the credit card market is commensurate with the market's importance for many households. Overall the US credit card market is large, with over $700B in outstanding balances, and is an important source of credit. Roughly 40% of US households borrow, or "revolve" a balance, on a credit card, and borrowing is especially prevalent among cardholders with lower credit scores. For example, among the lenders I study in this paper, over 80% of active accounts and 90% of account balances held by borrowers with FICO scores below 700 were revolving on an average day in 2008-2009 (see Table 2).

It is worth emphasizing that my study is focused solely on the CARD Act's pricing restrictions, and not on the Act's other sections that variously affected lenders' disclosures to borrowers, the timing of billing cycles, "nudges" for borrowers to pay more than their minimum payment, and more. This focus helps me to disentangle the repricing restrictions' effects from other regulation introduced in the same Act, and to shed light on a more general question of the optimal regulation of dynamic pricing in markets with information problems. In this sense, this paper is not an evaluation of the CARD

Act's effects overall per se (in the spirit of CFPB (2013) and Agarwal et al. (2015)), and more is a study of the equilibrium effects of a particular dynamic pricing regulation.

The remainder of this paper is organized as follows. Section II introduces the datasets used in the analysis, and provides background on the CARD Act. Section III presents descriptive evidence on the CARD Act repricing restrictions' effects on both lender rents and the pricing of risk. Section IV then develops my model to study these effects in equilibrium, and Section V presents progress toward estimating the model.

II. Data and Policy Background

This paper uses two large, anonymized datasets developed by the Consumer Financial Protection Bureau (CFPB). The first is the Credit Card Database (CCDB), a near-universe of de-identified US consumer credit card accounts, which I at times will refer to as the "account-level" data. The second is the Consumer Credit Panel (CCP), a large representative sample of de-identified US consumer credit reports from a national consumer credit reporting agency, which I at times will refer to as the "borrower-level" data.

The account-level CCDB and borrower-level CCP data cannot be linked, and each dataset provides different insight into the credit card market. On the one hand, the CCDB contains rich information on account pricing and usage, including the total amount spent, repaid, and borrowed each month on each account, and detailed price data, including fees of various types, total interest charges, the allocation of balances across various interest rates (e.g., a promotional rate vs. a standard rate). These data are available for the universe of credit card accounts held by over 20 large lenders under the supervisory authority of either the OCC or the CFPB, comprising over 85% of total outstanding US credit card balances. For much of this analysis, for data consistency reasons I use a subset of CCDB lenders; together these lenders hold nearly 80% of all credit card accounts observed in the CCDB in 2008-2009.

Meanwhile, the borrower-level CCP data provide rich information on how borrowers allocate balances across multiple credit cards and other loans, and how borrower behavior evolves over time across multiple accounts. In contrast with the CCDB, there is limited data on account-level pricing and usage. For example, account balances cannot be disaggregated into new spending, repayments, and revolved balances, and only for a select subset of credit card accounts is it possible to infer a measure of pricing, by backing out total fees and finance charges from the account's reported minimum payment. Also in contrast with the CCDB, the CCP is drawn from a 1/48 nationally representative random sample, rather than a near-universe.

These and similar data have been used in earlier academic research. The CCP data are utilized in Brevoort et al. (2016) and Brevoort and Kambara (2015), while the CCDB data are utilized in Alexandrov et al. (2016), Alexandrov and Grodzicki (2016), Keys and Wang (2016), and Gross et al. (2016). The CCDB data are a superset of the credit card account data studied in Agarwal et al. (2015) and Agarwal et al. (2016). The CCP data are similar to the Federal Reserve Bank of New York's consumer credit panel (FRBNY CCP), used in a number of studies including Brown et al. (2016); however, the CFPB CCP data show account-level information, rather than account-type aggregates as in the FRBNY CCP (for example, the balance on each of a borrower's credit cards, rather than the total balance summed across all credit cards).

Both datasets provide panel data on the credit card market spanning both before and after the CARD Act. The CCDB data are available from 2008 onward, and the CCP data begin still earlier, in 2001. Meanwhile the CARD Act was implemented in 2010, after a legislative and regulatory process that arguably began in earnest in December 2008, in the form of the Federal Reserve's final proposed revisions to its Unfair or Deceptive Acts or Practices regulation (UDAP, or Regulation AA). Many of the proposed regulations were incorporated into drafts of the CARD Act, which also added fee restrictions[2] and further limited the circumstances in which lenders could apply interest rate increases.[3] A full summary of the Act's restrictions and subsequent implementing regulation is available in CFPB (2013). The CARD Act was then signed in May 2009, and took effect in February 2010 and August 2010. This timeline is summarized in Figure 1 below.

Figure 1: CARD Act Timeline

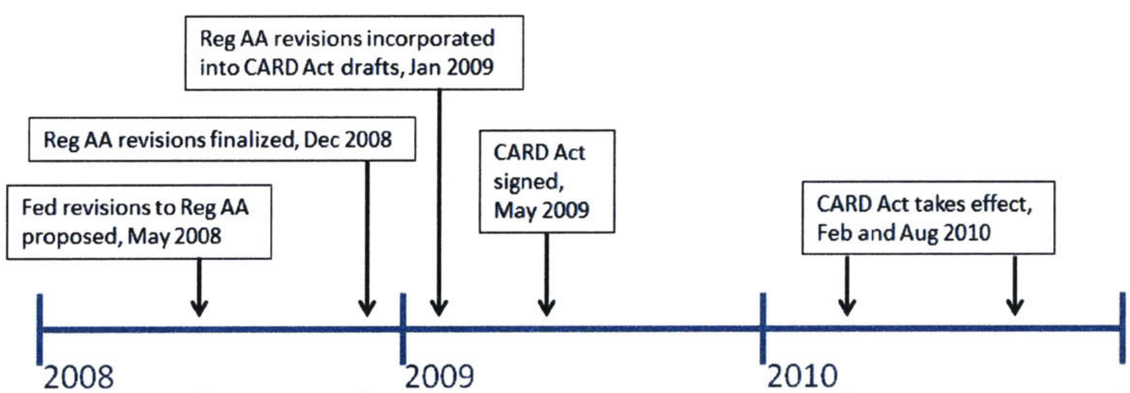

[2] Formally, these fee restrictions are implemented as requirement that fees be "reasonable and proportional," accompanied by a safe harbor regulation that late payment fees of $25 (or $35 for subsequent late payments within six months) are presumed to meet the reasonable and proportional requirement. Research by the CFPB (2013) indicates that "nearly every" lender complied with the safe harbor fee levels and that the fee levels were binding.

[3] The original Regulation AA revisions proposed to allow interest rate increases after late payments of 30 days or more; the CARD Act extended this to 60 days. Interestingly, the policy debate surrounding the Federal Reserve's original choice of 30 days emphasized that lenders would not benefit from rate increases at 60 days, given these accounts already-high loss rates (Federal Register 2009); this has borne out in practice given the near-zero use of 60-day-late interest rate increases among lenders.

III. Descriptive Results

In this section I present descriptive evidence that the CARD Act's pricing restrictions both lowered lender rents and also exacerbated the mispricing of risk. I then show evidence of how price changes affected borrower choices to participate in borrowing, and I discuss how these effects can interact to influence the overall efficiency of the credit card market.

My analysis of lender rents focuses on observed interest rate changes immediately after a borrower behavior that is *not* predictive of risk: late payment of a credit card bill by less than 30 days. In Table 1, I present coefficient estimates of β in the OLS model,

$$y_{it} = \alpha_0 + \alpha_{x(i,t)} + \alpha_{x_0(i)} + \beta \mathbf{1}_{\{\text{paid late} < \text{ 30 days}\}(i,t)} + \epsilon_{it}$$

where y_{it} represents the dependent variable indicated in each row of the table, and the α_x terms represent fixed effects for 10-point bins of borrowers' contemporaneous FICO score and origination FICO score. Data are at a monthly frequency. As shown in the first row of the table, lenders responded to late payments of less than 30 days by increasing interest rates by an average of 93 basis points in the subsequent month, in addition to any late fee charged. The second row then confirms that these price increases occurred despite these late payments not being predictive of risk, as measured by subsequent default.[4] In fact, the coefficient estimate in the second row suggests that borrowers paying late by less than 30 days are marginally *less* likely to default in the following year than their peers, by roughly 0.1 percentage points, relative to a sample mean of 13 percentage points.

It is worth emphasizing that these late payments are not reported to credit bureaus and therefore are private information for the borrower's lenders. So, these subsequent price increases are

[4] My measure of subsequent default is the one most relevant for lender profitability in the medium run – a loan being charged off (written off as a loss) at any point during or before the year after the sample period (i.e. during or before 2010 for the pre-CARD-Act sample, and during or before 2013 for the post-CARD-Act sample).

unlikely to reflect changes in borrowers' pricing offers from competing lenders, and rather reflect rents earned without an accompanying increase in borrower risk. Reasons why lenders raised prices on these borrowers are difficult to test directly in the data, but casual empiricism would suggest that the same borrowers who tend to miss a repayment due date by a few days – perhaps these are busy borrowers, "disorganized" borrowers, or borrowers with short-duration liquidity shortfalls (Leary and Wang, 2016) – may also be borrowers who are less likely to have the time or wherewithal to pay off their balance or transfer their balance to another lender.

I then confirm in the second column of Table 1 that the CARD Act effectively eliminated such repricing after a "barely" late payment. The coefficient estimates indicate that the average subsequent interest rate change was nearly zero (1 basis point) for a "barely" late payment incurred after the CARD Act, whereas such late payments predict higher default rates in the post-CARD-Act sample period by only 0.002%.

While the preceding analysis is indicative of lenders using repricing to earn rents on private information about borrower demand, there is also evidence that lenders possessed, and priced on private information about borrower risk. To show this, I ask how ex-post risk (default) outcomes differ across different price levels *after* controlling for all observable information about risk. In the present analysis, I focus only on 20-point bins of contemporaneous FICO score as my measure of observable risk, and I ask how residual default risk correlates with residual price levels. As shown in Figure 2, the relationship between the two is strong. The plot shows average ex-post default rates at vigintiles (20ths) of the distribution of fee-inclusive prices. Across the distribution and well into its tails, residual default risk is positively correlated with residual prices, suggesting that lender pricing strategies responded to private information about risk. While many sources of lender private information are

Table 1: Late Payments of Less than 30 Days

Dependent Variable	Pre-CARD Act Sample (2008-2009) Mean of Dep. Var.	Pre-CARD Act Sample (2008-2009) Paid Late by < 30 Days	Post-CARD Act Sample (2011-2012) Mean of Dep. Var.	Post-CARD Act Sample (2011-2012) Paid Late by < 30 Days
$\Delta_{t-1:t}$(APR)	0.255	0.927***	0.061	0.0130***
		(0.00237)		(0.00175)
Subsequent Default	0.130	-0.00119***	0.072	0.00277***
		(0.000306)		(0.000244)

Notes: Table shows sample means and regression coefficents for the dependent variable in each row. Regressors are a month-account level indicator for a borrower's repayment being late by less than 30 days. Regressions include controls for origination FICO and contemporaneous FICO in 10-point bins.

Figure 2: Delinquency Rates by Price Level, Net of Observable Risk

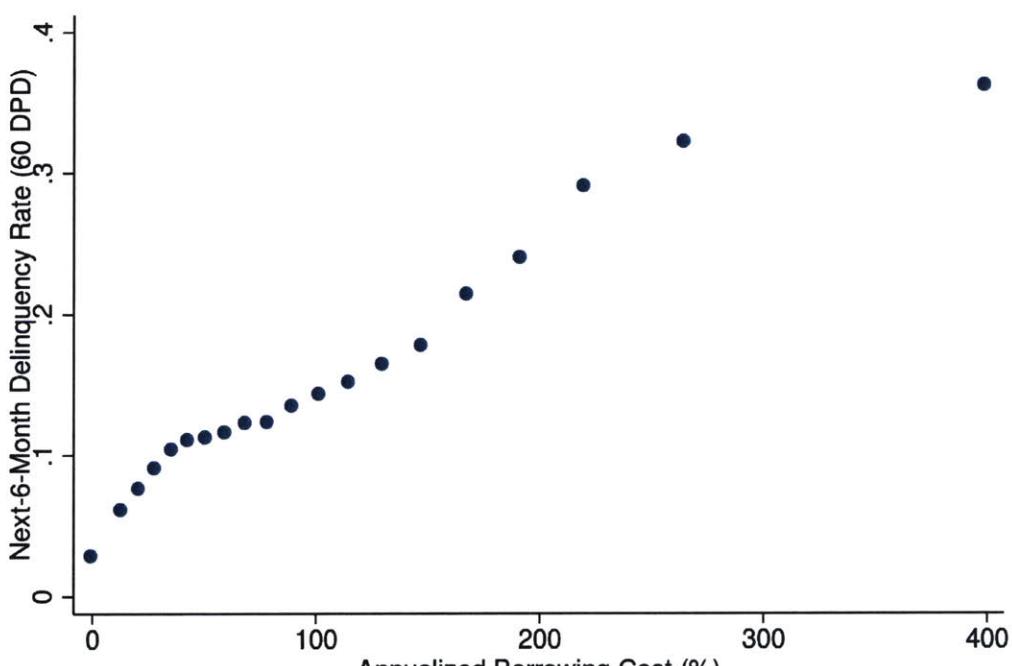

Note: Plot shows 60-day-past-due delinquency rates at a six-month horizon for quantiles of annualized (fee-inclusive) borrowing cost, after partialling out 20-point bins of borrower's contemporaneous FICO score.

unobservable for a researcher in the available data, such information sources abound for lenders: these include purchase behaviors, self-reported income changes, revolving status and its persistence, repayment habits (e.g. only paying the minimum amount due), and lagged values of balances and credit limits.

Just as lender pricing of private information about demand changed with the implementation of the CARD Act, so also did lender pricing of risk. I make this observation in a pair of figures (Figure 3A and Figure 3B), which respectively show data from pre-CARD-Act and post-CARD-Act periods. In each figure I plot two price gradients, one gradient showing the relationship between price and risk at the time of account origination, and the other gradient showing the relationship between price and newly realized risk on mature, already-originated accounts.[5] Notably the CARD Act pricing restrictions directly affected the pricing of newly realized risk, but not the pricing of risk at origination. Hence these figures present a "graphical difference in differences" of the price of newly realized risk relative to origination risk, after the Act relative to before.

Specifically, the two price gradients in each figure are generated as follows. For newly originated accounts, the gradient is straightforwardly the conditional means of account interest rates at quantiles of origination FICO scores; these gradients are plotted using hollow circles and dashed lines relative to the bottom, left axes. Both before and after the CARD Act, the lines of best fit for these conditional means suggest a price of risk at the time of origination of roughly 30 basis points for every 10 points of FICO score.

---

[5] In particular I focus on accounts originated at least 24 months prior, but no more than 60 months prior, so as to have an approximately equal distribution of account ages in both the pre- and post-CARD-Act samples.

Figure 3A: Pre-CARD-Act Price of Origination Risk and New Risk

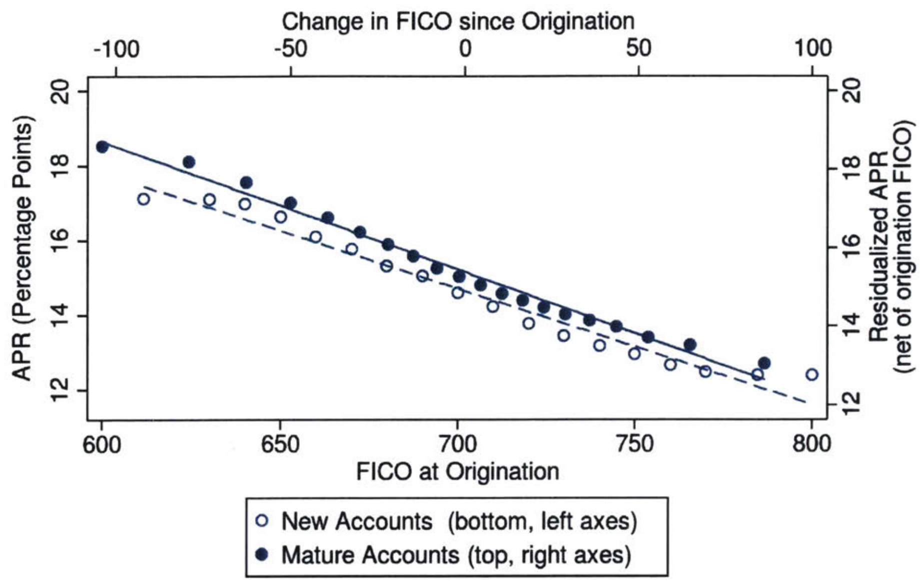

Notes: The post-CARD-Act sample is accounts originated after 2010, and observed at ages 24 months or more. The pre-CARD-Act sample is accounts observed prior to 2010 at ages 24 months or more, and at ages no greater than the maximum age observed in the post-CARD-Act sample. Residualization of APR is performed using dummies for 10-point FICO score bins, with sample averages added back in after residualization.

Figure 3B: Post-CARD-Act Price of Origination Risk and New Risk

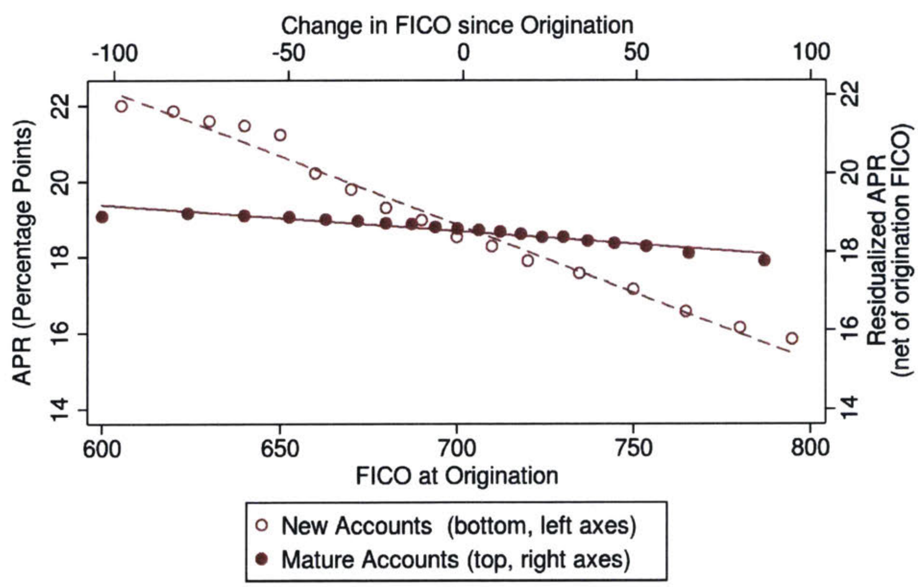

Notes: The post-CARD-Act sample is accounts originated after 2010, and observed at ages 24 months or more. The pre-CARD-Act sample is accounts observed prior to 2010 at ages 24 months or more, and at ages no greater than the maximum age observed in the post-CARD-Act sample. Residualization of APR is performed using dummies for 10-point FICO score bins, with sample averages added back in after residualization.

Meanwhile for mature, existing accounts, the displayed price gradient is only slightly more nuanced. Ideally the gradient would be the relationship between price changes since origination, $p_{it} - p_{i,0}$, and changes in FICO score since origination, $x_{it} - x_{i,0}$. However, prices at origination $p_{i,0}$ are unfortunately not observed for accounts originated prior to 2008, which includes all mature accounts in the pre-CARD-Act sample. Hence I rearrange the ideal regression,

$$p_{it} - p_{i,0} = \alpha_0 + \beta(x_{it} - x_{i,0}) + \epsilon_{it}$$

by moving the unobserved $p_{i,0}$ to the right-hand side and proxying for $p_{i,0}$ with controls for origination FICO score,[6]

$$p_{it} = \alpha_{x_{i,0}} + \beta(x_{it} - x_{i,0}) + \epsilon_{it}$$

The displayed gradients are then conditional means of p_{it} at quantiles of the regressor $(x_{it} - x_{i,0})$, after partialling out the fixed effects $\alpha_{x_{i,0}}$. So, the price gradient on new risk is interpretable as the price per point change in FICO score since origination. This gradient is plotted relative to the upper, right axes in both figures, on an identical scale as the opposite axes.

Turning to the figures, the striking result from both plots is that, before the CARD Act (Figure 3A), prices on mature accounts adjusted in response to changes in borrower FICO score, and in fact the price gradient of newly emergent risk was approximately the same as the price gradient of risk at origination; in contrast, after the CARD Act (Figure 3B), these two become decoupled, and in fact price is nearly unresponsive (flat) with respect to risk changes after origination. This underscores the effects

[6] Estimates of this proxy regression will only differ from coefficients that would be estimated in the ideal regression to the extent that the residual part of initial prices unexplained by initial FICO score, $\eta_i = p_{i,0} - \alpha_{x_{i,0}}$, covaries with the term $(x_{it} - x_{i,0})$. Also by Frisch-Waugh, the proxy regression is closely related to the simple regression of p_{it} on x_{it}, although these regressions differ insofar as $(x_{it} - x_{i,0})$ covaries with $x_{i,0}$, that is, the extent to which origination FICO predicts subsequent change in FICO.

of the CARD Act repricing restrictions: newly emergent risk changed from being priced nearly the same as origination risk, to being priced nearly not at all.

In order for this mispricing of risk to have an effect on the broader credit card market, it must affect borrowing choices, which I verify in the following figure. Figure 4 shows monthly hazard rates of exit from borrowing status as a function of change in FICO since origination, again plotted across quantiles of the regressor before and after the CARD Act. Across the distribution of FICO score changes, borrower exit rates evidently respond to the price incentives illustrated in the previous pair of figures: borrowers who became riskier over time, who benefited directly from the CARD Act pricing restrictions through lower borrowing costs on an existing account, become less likely to stop borrowing on that account; conversely borrowers who becomes safer over time become more likely to

Figure 4: Adverse Attrition by Changes in Risk

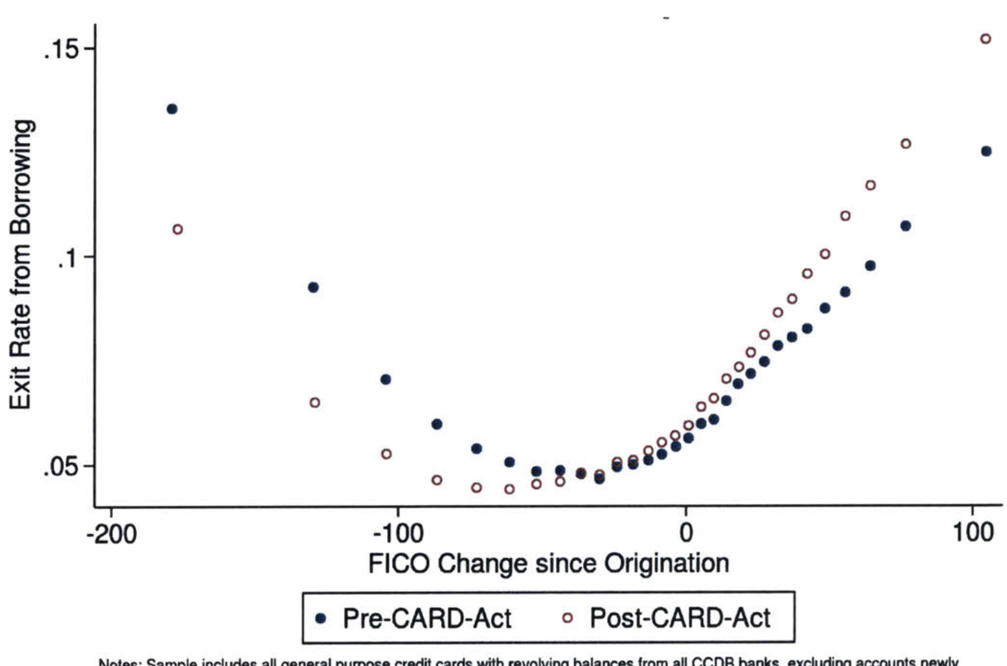

Notes: Sample includes all general purpose credit cards with revolving balances from all CCDB banks, excluding accounts newly opened in the prior 12 months.

stop borrowing on that account. Thus lenders became more likely to retain newly bad risks, and to face attrition of newly safe borrowers. These differential attrition incentives introduce an adverse selection problem – or adverse attrition problem – that as usual in equilibrium can drive up prices for all borrowers.

Before developing a framework to help evaluate these equilibrium effects, I conclude this section with a visual summary of some of the pricing changes that followed after CARD Act implementation. In Figure 5, I present the interquartile range (IQR) of interest rates charged on mature credit card accounts over time, after controlling for observable borrower risk (here taken to be 20-point bins of contemporaneous borrower FICO score). The data are shown separately by quarterly origination vintage, so that each line represents the IQR over time for accounts originated in a given quarter. Lines that end earlier indicate earlier origination vintages, as each vintage is tracked for exactly five years.

The striking feature on the plot is the sharp and persistent drop in interest rate IQRs that coincides with the implementation date of the CARD Act's interest rate repricing restrictions, denoted by the vertical black line. Before the CARD Act, the average IQR of interest rates (controlling for risk) was roughly 8% on average across account vintages; immediately after the Act's implementation this fell sharply to less than 6% on subsequent vintages of mature accounts. Furthermore, this drop in price dispersion arises despite the effects seen previously in Figure 3B, which necessarily imply higher price dispersion net of observable risk.

As a benchmark, the pre-CARD-Act IQRs of 8% shown in the figure are roughly equal to the difference in the CCDB data in total (fee-inclusive) average borrowing cost between a FICO-700 borrower and a FICO-800 borrower. Hence, any analysis of the CARD Act's pricing effects that focuses only on averages is likely to miss important distributional effects: before the Act, prices differed sharply across borrowers within identical risk groups, with, as I have argued in this section, some of that price dispersion arising because of lender pricing to borrower demand, and some arising because of lender private information about risk. The model in the following section helps quantify how much the observed equilibrium drop in price dispersion reflects lower lender rents, versus weaker pricing of risk, and how these changes affected the overall efficiency of the market.

Figure 5: Price Dispersion on Mature Accounts by Vintage

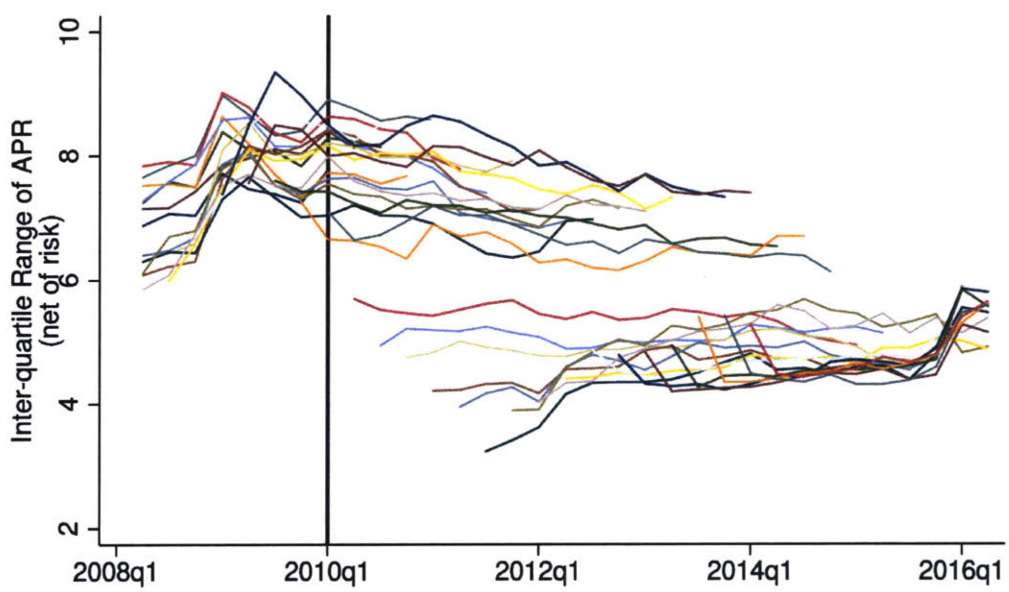

Notes: Plot shows the interquartile range of APR residuals after partialling out 10-point bins of origination FICO and contemporaneous FICO. Each line represents a different vintage of credit card originations, at the quarterly level. Sample includes all mature, revolving general-purpose credit card accounts from in-sample CCDB banks. Mature accounts are defined by account ages of 18+ months since origination, so as to remove initial 'teaser' APRs from the sample. Vertical black bar shows the implementation date of CARD Act APR re-pricing restrictions.

IV. Model

To study the CARD Act's price restriction effects in equilibrium, I build a model of the credit card market including the key features suggested by the analysis so far: heterogeneous credit demand

and risk across consumers, public and private information about consumer types, consumer types that change over time, and imperfectly competitive lenders who earn rents from at least some of their borrowers. At the heart of the model are differentiated lenders competing for borrowers who choose in each period whether to borrow on a credit card, whether to hold a credit card without borrowing, and whether to switch their cardholding from one lender to another. In my exposition of the model, I begin by discussing primitives on the consumer side – how consumer types differ and how types evolve – and then turn to discussing supply primitives.

To keep the model tractable, I restrict consumer type heterogeneity to consist of one dimension of publicly observable information – understood here as a FICO score – and one dimension of privately observable information. I denote the former as $x_{i,t}$ and the latter as ψ_{it} for borrower i in period t, sometimes suppressing subscripts for brevity. Each consumer's private-information type ψ therefore encodes any difference in both risk and credit demand across consumers that is not predicted by FICO score. Jointly these two dimensions determine borrower types, which again for brevity I will denote as $\theta = (x, \psi)$.

Time is discrete. At the start of each period, consumers exogenously receive a new type θ. Consumers observe their types and then choose whether to (1) borrow on their credit card, or (2) hold a credit card without borrowing on it, or (3) exit from the credit card market altogether. Consumers who do not exit from the market also choose whether to remain with their current lender, or to switch to a new lender. The number of consumer choices is therefore 2J+1, where J is the number of lenders.

After making these choices, consumers enjoy flow utilities as follows. Consumers who choose to borrow enjoy a flow utility of borrowing from lender j, $D = D(\theta, j)$. Borrowers who change lenders incur a switch cost, $S = S(\theta)$. Borrowers also incur disutility $-\gamma p$ when borrowing at price p, if they do not default on their loan, which occurs with probability $\delta(\theta)$. Finally, to help rationalize different markets shares across lenders in the data, each lender j provides consumers with a flow utility $\xi_j$regardless of whether those consumers borrow, or simply hold the lender's credit card without borrowing. Each choice k also provides consumers an i.i.d., extreme-value type-one distributed taste shock, so that consumers have logit demand over this choice set. For the outside good – that is, exiting the credit card market – consumer utility is as usual normalized to zero net of these taste shocks.

To summarize, expected consumer flow utility from choice k can be written as follows, where $j(k)$ is used to denote the lender chosen in choice k, and $j(i,t-1)$ is the consumer's preceding-period lender, and $1_{D(k)}$ is an indicator for whether k is a choice to borrow from, rather than simply hold a credit card from, lender $j(k)$:

$$U_{ik} = 1_{D(k)} D\big(\theta, j(k)\big) - 1_{\{j(k) \neq j(i,t-1)\}} S(\theta) - (1 - \delta(\theta)) \gamma p_{j(k)} + \epsilon_{ik}$$

Because these flow utilities differ across borrower types θ, the demand system can be understood as a flexible form of random coefficients logit, with the random coefficients governed by the terms $D\big(\theta, j(k)\big)$. As in Fox et al. (2016), this approach is nonparametric about the distribution of the random coefficients, while approximating the overall distribution of consumer types with a (potentially large) finite number of types θ.

Given that consumers are solving a dynamic problem in choosing their lender, the total expected utility from choice k also depends on a discounted, expected continuation value $\beta \mathrm{E}[V(\theta', j(k))]$, where θ' is the borrower's next-period type. Borrowers who default (again, with exogenous probability $\delta(\theta)$) are assigned to have no lender, i.e. the outside good, at the end of the period, but are allowed to borrow again in subsequent periods, so that defaulting borrowers likewise have expected continuation value $\mathrm{E}[V(\theta', 0)]$.

Borrower types are assumed to evolve according to a first order Markov process, with transition matrix denoted T_θ. The Bellman equation for borrower continuation values can therefore be tractably written in vectorized form as,

$$V(\theta,\hat{j}) = \max_k \left[1_{D(k)}D\big(\theta,j(k)\big) - 1_{\{j(k)\neq \hat{j}\}}S(\theta) - \big(1-\delta(\theta)\big)\gamma p_{j(k)} + \epsilon_{ik}\right.$$

$$\left.+\beta\big(1-\delta(\theta)\big)T_{\theta:\theta'}V(\theta',\hat{j}(k))\right] + \beta\delta(\theta)T_{\theta:\theta'}V(\theta',0)$$

On the supply side, at the start of each period, lenders post "teaser" prices to attract new borrowers and also post a separate set of prices for returning borrowers. A crucial assumption in the model, building on the mechanism in Sharpe (1990), is that lenders fully observe consumer types $\theta = (x,\psi)$ for returning consumers, whereas lenders only observe publicly observable types x for new consumers. Hence teaser rates can be written as $p_0(j,x)$ for lender j, and returning borrowers' prices can be written as $p_1(j,x,\psi)$ or, for brevity, $p_1(j,\theta)$.

Lenders are differentiated insofar as, for two different lenders j and j', unobserved qualities ξ_j and $\xi_{j'}$ differ and consumer-specific borrowing utilities $D(\theta,j)$ and $D(\theta,j')$ differ. This specification of lender differentiation is agnostic as to whether horizontal or vertical differentiation is dominant in the market, and as to the relative importance of differentiation vis-à-vis other demand determinants such as price. This specification also allows some lenders to have specialization or competitive advantages in some parts of the lending market, for example, by providing higher borrowing utility to consumers of some types θ. This accords with evidence both about how lenders differentiate their product offers (Ru and Schoar, 2016) and evidence on how lender market shares differ substantially across FICO score groups.

Besides lender differentiation, lender costs are the other primary force driving lender pricing. Lenders incur a marginal per-period cost $c(\theta,j)$ for consumers who choose to borrow, and additionally incur a one-time acquisition cost $\kappa(x,j)$ for newly originated accounts. The separate specification of acquisition costs reflects the high price of customer acquisition in the credit card market, which relies both on new-customer incentives and voluminous direct-mail credit card offers (Grodzicki, 2015).

For simplicity, these two types of costs are assumed to be summary measures of all costs incurred regardless of account default or delinquency.[7] Nevertheless, it is important to emphasize that some of these costs may, on net, be negative, as there are important sources of lender revenue – namely, interchange revenue off of transactional use of credit cards, and cross-sell opportunities to borrowers for whom a credit-card may be a "gateway" credit product – that are not elsewhere included in the model.

In summary, the model includes a rich set of important features in the credit market and also has some notable limitations. The model focuses only on the extensive margin of borrowing, leaving the intensive margin, i.e. borrowers' choice of how much to borrow, for later work. The model also assumes borrowers "single-home," i.e. hold only on credit card at a time, or equivalently make multiple, independent borrowing choices across their multiple cards. The model also does not differentiate between fees and interest charges, instead combining both into a single annualized "price" normalized by the credit limit on the card.[8]

---

[7] Hence these costs can equivalently be understood as per-period insurance payments against borrower default and all other idiosyncratic account-specific costs.

[8] This choice of normalization is one of several options – a leading alternative would be to normalize by a consumer's balance on the card. The difficulty with normalizing by balances is it generates severe period-to-period variability in borrowing costs for the same borrower, especially on low credit-limit cards. For example, a $500 credit limit card carrying a $100 balance in one period and a $200 balance in the next, paying $0 in interest and $25 in fees, would have a period-to-period change in annualized borrowing cost from 75% APR-equivalent to 150% APR-equivalent, conceivably without any actual change in the borrower's risk from the lender's perspective, or the cost of the loan from the borrower's perspective.

V. Model Estimates

This section presents preliminary progress toward estimating the model. To take the model to the data, the first step empirically is to define borrower types. Starting with publicly observable types, I recognize, as in Table 2, that borrowing rates fall to substantially low levels at FICO scores of 800 and above (fewer than 25% of active accounts carry a balance). I therefore pool together borrowers of FICO score 780 and above as "high-FICO" types. For all other borrowers, I group borrowers by 20-point FICO bins. These bins are the set X from which observable types x are drawn.

Table 2: Revolving Rates by FICO Score Band

		Revolving Rates	
		Share of Active Accounts	Share of Balances
All Cardholders		65.03%	84.46%
FICO Score	< 600	93.56%	97.15%
	600 - 620	89.75%	95.86%
	620 - 640	88.99%	95.43%
	640 - 660	87.79%	94.84%
	660 - 680	85.04%	93.64%
	680 - 700	80.78%	92.08%
	700 - 720	75.50%	90.01%
	720 - 740	69.94%	87.41%
	740 - 760	62.55%	82.58%
	760 - 780	50.14%	72.74%
	780 - 800	34.27%	56.52%
	800 - 820	24.00%	41.49%
	820 - 840	21.63%	38.61%
	840+	18.60%	33.40%

Notes: Table shows share of active accounts (column 1) and total balances (column 2) on which borrowers do not repay their balance in full, and thus pay to borrow ("revolve"), for each FICO score range. Pooled rates are reported for all in-sample lenders.

The more difficult part of defining borrower types is to identify private-information types. My approach here is motivated by the pattern seen previously in Figure 2, namely that *within* observable risk types, unobservable risk (as measured by subsequent risk outcomes) is increasing in prices charged by lenders. I therefore plot, for each lender in my sample and for each observable risk type, isotonic fits of lender pricing functions for borrowers with each default probability. Formally, this is accomplished through an isotonic linear probability regression of realized default outcomes on price, followed by inversion of the resulting isotonic fits.[9]

In principle, this procedure recovers a continuum of borrower types (across a compact set of observed default probabilities), but for tractability I take averages and aggregate these types into discrete private-information types ψ. In my baseline specification, I allow for 10 different private-information types at each publicly observable type (FICO score) x. Each public-private pair, $\theta = (x, \psi)$ then has a unique default probability that is common across lenders, and a returning-borrower price $p(\theta, j)$ that is distinct across lenders j. After this process is repeated for each FICO group, the result is a set of pricing functions across borrower types.

After estimating these private information types, it is straightforward to estimate the joint transition probabilities of public and private types, the aforementioned Markov transition matrix T_θ, directly in the data. Transition probabilities are estimated only on consumers who borrow from the same lender for two consecutive periods, because by the exogeneity of T_θ this approach introduces no bias in the resulting estimates of T_θ.

The estimated transition matrix is shown as a contour plot in Figure 6 below. For ease of labeling the axes, borrower types are written in a single index where the units digit corresponds

[9] This approach partly draws on the motivation for isotonic regression in Hausman et al. (1998), in cases where theory predicts an isotone relationship but finite data present non-isotonic exceptions.

observable types x, and the tenths digit corresponds to private-information types ψ. So for example, borrower type 1.9 is in the lowest FICO score bin, and the highest private-information category (corresponding to the highest default rates within that FICO score bin).

As can be seen, there is strong persistence along the diagonal, and most transitions are within no more than 40 points' change in FICO score at a quarterly frequency. For borrowers who do change FICO group, private-information types are still somewhat persistent, which generates the ripple pattern seen off the graph's primary diagonal. That is to say, borrowers who have high expected default rates among, for example, 680-FICO borrowers, are likely to persist in having high default rates relative to their peers even after transitioning to being, for example, a 660-FICO borrower.

Figure 6: Borrower-Type Transition Probabilities

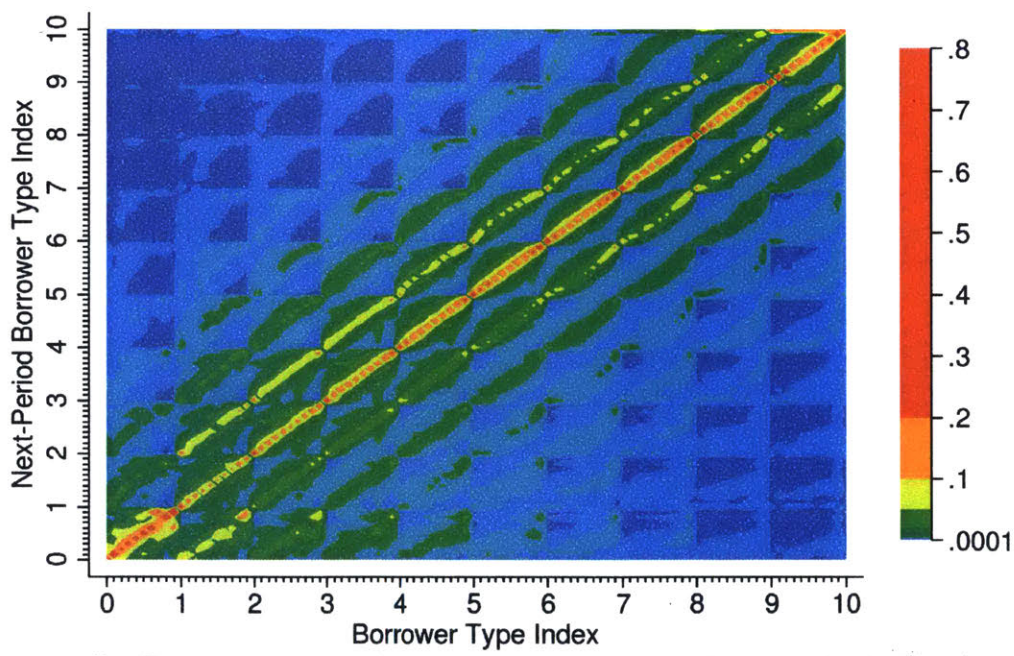

Notes: Plot shows quarterly transition probabilities among borrower types. Borrower type indexes are partitioned by whole numbers corresponding to 20-point FICO bins ranging from 580 to 780, and subdivisions that index private-information types. For example, indexes between 0 and 1 cover the range of private-information types among FICO-580 borrowers, indexes between 1 and 2 cover the range of private-information types among FICO-600 borrowers, etc.

A next step in estimating the model is to measure borrowers' price sensitivity, captured in the model by the marginal disutility of price, γ. As is well known (Berry, 1994), estimating price sensitivity in an imperfectly competitive market requires exogenous variation in prices in order to separately identify price effects from the effects of firm quality, such as the lender qualities ξ_j here. To estimate γ, I leverage a novel source of price variation in consumer credit markets – idiosyncratic, portfolio-wide repricing by certain lenders – so as to ground the model in credible demand estimates.

Figure 7 below illustrates an example of such a repricing campaign. This campaign increased nearly all of a given lender's mature revolving accounts' interest rates by 100 basis points in the period labeled event-time 0, and subsequent borrower attrition from borrowing is shown in the second panel. This change in price and change in borrowing can be straightforwardly used to estimate a price

Figure 7: Portfolio-wide Repricing and Subsequent Attrition

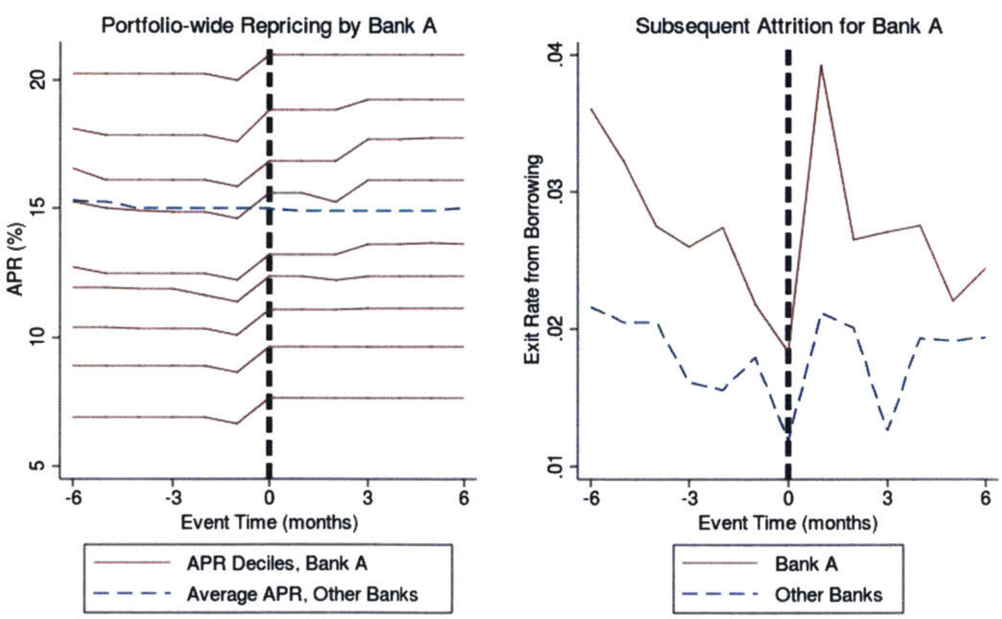

Notes: Sample includes all general purpose credit cards with revolving balances from all CCDB banks, excluding accounts newly opened in the prior 12 months.

elasticity of borrowing demand for each borrower type, which can then be translated into a marginal utility of income (or disutility from price). Specifically, let s_θ be the share of returning consumers of type θ choosing to borrow from the lender in question, and let p_θ be their price of borrowing, and let η_θ be these consumers' observed demand elasticity in the above experiment. Then leveraging the logit demand structure, the disutility of price γ satisfies,

$$\eta_\theta = -\gamma p_\theta (1 - s_\theta)$$

With multiple borrower types this disutility γ is over-identified, although for simplicity I allow a single value of γ that is common across all borrower types. In my analysis to date, my preferred estimate of γ is .173, as reported in Table 4 below.

With the preceding estimates in hand, estimates of the remaining model parameters are recovered using GMM, where model-determined likelihoods are made to match the following moments: (1) the probability of each borrower type on each lender continuing to borrow from that lender in the subsequent period; (2) the probability of each borrower type choosing instead to transact (carry a card, but not revolve a balance) from that lender in the subsequent period; and (3) the market share of each lender among all open accounts;. Additionally, the first-order conditions implied by lender profit maximization are used in the GMM procedure to identify lenders marginal costs for each retained account, and lender acquisition costs for each new account.

The remaining tables in this section present demand-side parameter estimate from this GMM procedure. Several patterns in the results are worthy of note. First, lender marginal costs are positively correlated with borrowing demand across the range of borrower types. This positive correlation between cost and willingness to pay is evidence of adverse selection on price: borrowers who are willing to pay higher prices are also more costly for lenders. Second, switching costs and borrowing demand are of comparable orders of magnitude, reinforcing the importance of including switch cost parameters in the model. Third, in-sample lenders are found to be substantially horizontally differentiated, with the most "subprime" of these lenders being preferred by borrowers with FICO scores of approximately 640, and the most "prime" of these lenders being preferred by borrowers with FICO scores of approximately 780. These results likewise point to the importance of allowing for lender differentiation in the model.

The next step in the analysis is to use these parameter estimates to calculate optimal lender prices in a model adapted to reflect CARD Act repricing restrictions. While this final step is nontrivial, it builds on the estimates attained so far. The optimal prices calculated in the adapted model can then shed insight on how various borrower types' prices change after the implementation of such price restrictions, whether these price changes on net bring the market closer to or farther from an efficient benchmark of prices equal to marginal costs, and how these changes affect the net subsidy paid across borrowers of different types and different profitability. For example, did the CARD Act repricing restrictions raise or lower the cross-subsidy paid to borrowers whose riskiness deteriorates rapidly? How much did the adverse selection effects shown previously in Figure 4 drive up borrowing costs across the market? Which types of borrowers, both in the cross-section and as borrower types change over time, faced the greatest positive and negative price changes from the Act's pricing restrictions? These questions and more remain for future work.

Table 3: Demand Parameter Estimates, by Borrower Type

FICO Group	Borrower Type Index	Average Demand for Borrowing	Switch Cost	Default Probabiltiy	FICO Group	Borrower Type Index	Average Demand for Borrowing	Switch Cost	Default Probabiltiy
580	1	11.055	9.948	0.009	680	51	1.948	11.493	0.003
580	2	14.057	12.225	0.025	680	52	2.550	13.269	0.004
580	3	13.079	12.516	0.031	680	53	3.267	11.139	0.005
580	4	18.824	13.140	0.047	680	54	12.709	19.319	0.006
580	5	14.826	18.533	0.072	680	55	4.213	11.129	0.007
580	6	18.366	17.949	0.094	680	56	4.262	11.287	0.008
580	7	-17.409	-15.465	0.119	680	57	4.757	10.810	0.009
580	8	19.981	22.548	0.141	680	58	4.790	11.647	0.012
580	9	14.227	10.683	0.174	680	59	8.470	16.084	0.016
580	10	-18.844	-7.411	0.269	680	60	8.972	14.320	0.028
600	11	9.110	3.451	0.007	700	61	1.588	9.298	0.002
600	12	13.117	6.516	0.009	700	62	2.241	10.004	0.003
600	13	9.077	8.251	0.012	700	63	2.390	10.384	0.003
600	14	9.629	8.092	0.014	700	64	2.660	9.743	0.004
600	15	14.056	7.859	0.022	700	65	3.105	12.247	0.005
600	16	11.785	11.334	0.029	700	66	3.346	13.595	0.006
600	17	21.122	19.936	0.030	700	67	3.623	14.174	0.007
600	18	16.562	18.834	0.037	700	68	3.713	11.012	0.008
600	19	-9.953	-18.666	0.047	700	69	5.420	10.847	0.011
600	20	15.837	12.654	0.091	700	70	6.961	11.111	0.020
620	21	4.307	18.004	0.007	720	71	1.442	5.554	0.001
620	22	6.410	8.214	0.008	720	72	2.043	8.527	0.002
620	23	7.517	9.316	0.009	720	73	1.981	5.850	0.002
620	24	7.664	9.490	0.012	720	74	2.294	8.750	0.002
620	25	7.681	9.615	0.015	720	75	2.490	8.941	0.003
620	26	13.521	8.247	0.020	720	76	2.669	9.928	0.004
620	27	14.116	10.736	0.023	720	77	2.950	10.605	0.004
620	28	16.322	21.690	0.027	720	78	3.233	25.262	0.005
620	29	10.161	14.755	0.035	720	79	3.772	13.726	0.007
620	30	9.486	17.275	0.072	720	80	-17.798	-263.787	0.014
640	31	3.393	11.178	0.006	740	81	1.462	13.320	0.001
640	32	4.844	6.990	0.006	740	82	1.220	8.934	0.001
640	33	5.569	9.157	0.008	740	83	1.600	5.267	0.001
640	34	5.826	9.252	0.010	740	84	1.834	8.829	0.001
640	35	18.209	6.874	0.011	740	85	1.968	9.461	0.002
640	36	6.492	9.555	0.014	740	86	2.186	11.746	0.002
640	37	9.744	8.103	0.017	740	87	2.315	11.382	0.002
640	38	15.331	10.754	0.023	740	88	2.678	11.740	0.003
640	39	13.486	18.303	0.028	740	89	-22.545	-90.016	0.004
640	40	-9.455	-89.706	0.044	740	90	-18.008	-129.003	0.008
660	41	2.333	14.191	0.004	760	91	0.882	8.648	0.000
660	42	17.050	9.192	0.006	760	92	0.804	9.427	0.001
660	43	4.527	7.886	0.006	760	93	1.208	9.205	0.001
660	44	5.056	8.229	0.008	760	94	1.335	9.377	0.001
660	45	5.492	9.492	0.009	760	95	1.341	10.371	0.001
660	46	5.642	10.595	0.010	760	96	1.538	9.655	0.001
660	47	6.251	10.860	0.012	760	97	1.559	11.875	0.001
660	48	8.579	8.764	0.017	760	98	-18.091	-54.177	0.002
660	49	7.594	13.710	0.021	760	99	-18.324	-56.631	0.002
660	50	11.653	15.266	0.035	760	100	-18.170	-66.294	0.005
					780	101	-34.788	-0.425	0.001

Table 4: Demand Parameter Estimates, Common Across Types

	Parameter Estimate
Implied Hotelling Travel Cost, Public Types	4.892
Implied Hotelling Travel Cost, Private Types	3.136
Marginal Utility of Income (Disutility of Price)	0.1726

WORKS CITED

Agarwal, Sumit, Souphala Chomsisengphet, and Chunlin Liu. "The importance of adverse selection in the credit card market: Evidence from randomized trials of credit card solicitations." Journal of Money, Credit and Banking 42, no. 4 (2010): 743-754.

Agarwal, Sumit, Souphala Chomsisengphet, Neale Mahoney, and Johannes Stroebel. "Do banks pass through credit expansions? the marginal profitability of consumer lending during the great recession." NBER Working Paper w21567 (2015).

Agarwal, Sumit, Souphala Chomsisengphet, Neale Mahoney, and Johannes Stroebel. "Regulating Consumer Financial Products: Evidence from Credit Cards." The Quarterly Journal of Economics 1 (2015): 54.

Alexandrov, Alexei and Daniel Grodzicki. "Credit Card Indebtedness and Community Influence." Working Paper, 2016.

Alexandrov, Alexei, Özlem Bedre-Defolie, and Daniel Grodzicki. "The Effects of Interest Rate Changes and Add-on Fee Regulation on Consumer Behavior in the U.S. Credit Card Market." Working Paper, 2016.

Ausubel, Lawrence M. "The failure of competition in the credit card market." The American Economic Review (1991): 50-81.

Ausubel, Lawrence M. Adverse selection in the credit card market. working paper, University of Maryland, 1999.

Bar-Gill, Oren, and Ryan Bubb. "Credit card pricing: The card act and beyond." Cornell L. Rev. 97 (2011): 967.

Berlin, Mitchell, and Loretta J. Mester. "Credit card rates and consumer search." Review of Financial Economics 13, no. 1 (2004): 179-198.

Berry, Steven T. "Estimating discrete-choice models of product differentiation." The RAND Journal of Economics (1994): 242-262.

Brevoort, Kenneth P., and Michelle Kambara. "Are All Collections Equal? The Case of Medical Debt." JOURNAL OF CREDIT RISK 11, no. 4 (2015): 73-97.

Brevoort, Kenneth P., Daniel Grodzicki, and Martin Hackmann. "Medicaid and Financial Health." Working Paper, 2016.

Brito, Dagobert L., and Peter R. Hartley. "Consumer rationality and credit cards." Journal of Political Economy (1995): 400-433.

Brown, Meta, John Grigsby, Wilbert van der Klaauw, Jaya Wen, and Basit Zafar. "Financial education and the debt behavior of the young." Review of Financial Studies (2016): hhw006.

Bureau, Consumer Financial Protection. "'CARD Act Report." (2013).

Calem, Paul S., and Loretta J. Mester. "Consumer behavior and the stickiness of credit-card interest rates." The American Economic Review 85, no. 5 (1995): 1327-1336.

Cochrane, John H. "Time-consistent health insurance." Journal of Political Economy (1995): 445-473.

Drozd, Lukasz A., and Ricardo Serrano-Padial. "Modeling the Credit Card Revolution: The Role of Debt Collection and Informal Bankruptcy." (2013).

Drozd, Lukasz A., and Ricardo Serrano-Padial. "On Prepayment & Rollover Risk in the US Credit Card Market." (2014).

Finkelstein, Amy, Kathleen McGarry, and Amir Sufi. "Dynamic Inefficiencies in Insurance Markets: Evidence from Long-Term Care Insurance." American Economic Review 95, no. 2 (2005): 224-228.

Fox, Jeremy T., Kyoo il Kim, Stephen P. Ryan, and Patrick Bajari. "The random coefficients logit model is identified." Journal of Econometrics 166, no. 2 (2012): 204-212.

Fulford, Scott L. "How important is variability in consumer credit limits?." Journal of Monetary Economics 72 (2015): 42-63.

Ganong, Peter, and Pascal Noel. How Does Unemployment Affect Consumer Spending?. Working paper, 2015.

Grodzicki, Daniel. "Competition and Customer Acquisition in the US Credit Card Market." (2014).

Grodzicki, Daniel. The evolution of competition in the credit card market. Working paper. Stanford University, Department of Economics, Palo Alto, Calif, 2012.

Gross, David B., and Nicholas S. Souleles. "Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence from Credit Card Data." Quarterly Journal of Economics 117, no. 1 (2002).

Gross, Tal, Matthew J. Notowidigdo, and Jialan Wang. The marginal propensity to consume over the business cycle. No. w22518. National Bureau of Economic Research, 2016.

Handel, Ben, Igal Hendel, and Michael D. Whinston. "Equilibria in Health Exchanges: Adverse Selection versus Reclassification Risk." Econometrica 83, no. 4 (2015): 1261-1313.

Handel, Benjamin R. "Adverse selection and inertia in health insurance markets: When nudging hurts." The American Economic Review 103, no. 7 (2013): 2643-2682.

Hausman, Jerry A., Jason Abrevaya, and Fiona M. Scott-Morton. "Misclassification of the dependent variable in a discrete-response setting." Journal of Econometrics 87, no. 2 (1998): 239-269.

Heidhues, Paul, and Boton Koszegi. "Exploiting naivete about self-control in the credit market." The American Economic Review 100, no. 5 (2010): 2279-2303.

Hong, Suting, Robert M. Hunt, and Konstantinos Serfes. "Dynamic Pricing of Credit Cards and the Effects of Regulation." (2014).

Illanes, Gastón. "Essays on strategic behavior in government-designed markets." PhD diss., Massachusetts Institute of Technology, 2016.

Jambulapati, Vikram, and Joanna Stavins. "Credit CARD Act of 2009: What did banks do?." Journal of Banking & Finance 46 (2014): 21-30.

Keys, Benjamin and Jialan Wang. "Minimum Payments and Debt Paydown in Consumer Credit Cards." NBER Working Paper w22742 (2016).

Leary, Jesse B., Consumer Financial Protection Bureau, and Jialan Wang. "Liquidity Constraints and Budgeting Mistakes: Evidence from Social Security Recipients."

Levitin, Adam J. "Rate-Jacking: Risk-Based & Opportunistic Pricing in Credit Cards." Utah L. Rev. (2011): 339.

Petersen, Mitchell A., and Raghuram G. Rajan. "The effect of credit market competition on lending relationships." The Quarterly Journal of Economics (1995): 407-443.

Polyakova, Maria. "Regulation of Insurance with Adverse Selection and Switching Costs: Evidence from Medicare Part D." NBER Working Paper Series (2015): 21541.

Ronen, Joshua, and Tiago Pinheiro. "The unintended consequences of the credit card act." Journal of Law, Finance, and Accounting (2014).

Ru, Hong, and Antoinette Schoar. Do credit card companies screen for behavioral biases?. No. w22360. National Bureau of Economic Research, 2016.

Santucci, Larry. "A tale of two vintages: credit limit management before and after the CARD act and Great Recession." FRB of Philadelphia-Payment Cards Center Discussion Paper 15-01 (2015).

Sharpe, Steven A. "Asymmetric information, bank lending, and implicit contracts: A stylized model of customer relationships." The journal of finance 45, no. 4 (1990): 1069-1087.

Stango, Victor, and Jonathan Zinman. "Borrowing High versus Borrowing Higher: Price Dispersion and Shopping Behavior in the US Credit Card Market." Review of Financial Studies (2015): hhv072.

Stango, Victor, and Jonathan Zinman. "Fuzzy math, disclosure regulation, and market outcomes: Evidence from truth-in-lending reform." Review of Financial Studies 24, no. 2 (2011): 506-534.

Sullivan, James X. "Borrowing during unemployment unsecured debt as a safety net." Journal of human resources 43, no. 2 (2008): 383-412.

Telyukova, Irina A., and Randall Wright. "A model of money and credit, with application to the credit card debt puzzle." The Review of Economic Studies 75, no. 2 (2008): 629-647.